



From the Editor in Chief...

Air System Information Management

Robert E. Filman • *RIACS/NASA Ames Research Center* • filman@computer.org

I flew to Washington, DC, last week — a trip rich in distributed information management. Buying tickets, at the gate, in flight, landing, and at the baggage claim, myriad messages about my reservation, the weather, our flight plans, gates, bags, and so forth flew among a variety of travel agency, airline, and US Federal Aviation Administration (FAA) computers and personnel. By and large, each kind of information ran on a particular application, often specialized to its own data formats and communications network.

I went to Washington to attend an FAA meeting on System-Wide Information Management (SWIM) for the National Airspace System (NAS) (www.nas-architecture.faa.gov/Tutorials/NAS101.cfm). NAS (and its information infrastructure, SWIM) is an attempt to bring greater regularity, efficiency, and uniformity to the collection of stovepipe applications now used to manage air traffic. Current systems hold information about flight plans, flight trajectories, weather, air turbulence, current and forecast weather, radar summaries, hazardous condition warnings, airport and airspace capacity constraints, temporary flight restrictions, and so forth.

Information moving among these stovepipe systems is usually mediated by people (for example, air traffic controllers) or single-purpose applications. People, whose intelligence is critical for difficult tasks and unusual circumstances, are not as efficient as computers for tasks that can be automated. Better information sharing can lead to higher system capacity, more efficient utilization, and safer operations. Better information sharing through greater automation is possible, though not necessarily easy.

An Airspace Utopia

Ideally, when an airplane flying over the Rocky Mountains encounters turbulence and shifts its course, all “interested” parties could be made aware of the situation. For example, pilots of other airplanes on that flight path could be warned to avoid the turbulence; workers at the baggage ser-

vice in San Francisco, having a better idea of the real arrival times, could reschedule the luggage carousels; operators of connections for the passengers on that flight could make appropriate decisions about holding (or not holding) and routing flights; air traffic controllers could verify that the plane’s new trajectory doesn’t collide with some other craft; climate researchers on long-term turbulence patterns could add that incident to their database; and countless other future applications that might care about airspace information could patch into the information network.

What do we demand of a system moving that much flight information? Important “ilities” include

- efficiency (you can’t tell everyone everything),
- evolvability (you don’t know all the future applications),
- scalability (this is a big system, it will grow, and you don’t want it to be architecturally limited in capacity or vulnerable to a few points of failure),
- maintainability (when things don’t work, we need to be able to quickly find out why),
- reliability (for obvious reasons),
- quality of service (you need to get important information to its destinations quickly, deferring the unimportant — information about a hijacking takes priority over climate research) and
- security (information should be injected into and removed from the system only by appropriate people)..

The definition of appropriate in this last point includes not only keeping black hats from injecting spurious information into the system, but also the commercial concerns about sharing business information. An airline is willing to tell traffic control that a plane is late. It might even be willing to share that information with its own ground operations. However, the airline will often want to keep that information not only from its competitors, but

also from its own gate personnel. (This saves the clerk at the gate from having to conceal a lie when trying to convince you not to switch to another airline just because your flight won't take off for another three hours.)

Foundations of an Airspace Utopia

The desire to inform “all interested parties” argues for a publish-and-subscribe architecture. In publish-and-subscribe, *publishers* generate *events* in the course of program execution. In our airspace example, the flight over the Rockies would publish events announcing the turbulence and its new routing. *Subscribers* present to the system *subscriptions*: descriptions of the kinds of events that interest them. When the system detects events that match subscriptions, it notifies the subscribers with those events.

There is a lot of design space within

the scope of publish-and-subscribe. The first important decision relates to the structure of events. Possible structures for events include record-like, free-form, tree-structured (such as XML), and objects. Concomitant with the structure of events are their guaranteed properties. For example, we can demand that every event be signed, time-stamped, and marked for expiration. We might wish to distinguish among the information in the event envelope (for example, its generator), its content, and annotations about it (such as when or where it happened or its accuracy or precision).

Subscriptions talk about something and provide various operators for describing what makes events interesting. The simplest subscription mechanism is one based purely on the event's type — for example, a subscription to all turbulence events. While type-based subscriptions might work for controlling a workstation user interface, the distri-

bution and complexity of the airspace system demands a more refined approach. A better subscription language allows describing a predicate over the envelope, content, and annotations of a single event — for example, all events whose “type” is “turbulence,” whose location is within 300 miles of Denver, and which happened on 19 January.

The volume of information circulating in a national air system could demand a richer notion of subscription. We might find ourselves wanting to subscribe to a sequence of events that match some temporal property (for example, “Five or more reports of late takeoffs from an airport within an hour.”), that select from the set of generated events (for example, “The location of Flight 007 every three minutes.”), or even to embed programs to winnow through events within subscriptions (agent-oriented subscriptions). Subscriptions might express choreography

for the results or have their effects limited by the subscriber's access privileges. There is a tension between richness of expression in the subscription language and efficiency: we want to allow applications to precisely specify what information they want, but also want to minimize the work involved in culling that information for them.

Communication relies on shared structure and shared understanding. (It also relies on actual communication channels, the efficient maintenance of which is challenging for communicants moving at 1000 kilometers/hour. However, concrete communication is beyond the scope of this abstract discussion.) Part of developing such a system includes making decisions on common data formats (such as XML) and common data meanings (such as ontologies for airspace management). Thus, a subscription to turbulence events is likely to refer to an ontology of wind conditions. A clever system might even provide automatic or semiautomatic mechanisms for translating between representations, much as a client that generated

a subscription to aircraft location in terms of Pacific Standard Time might see an automatic transformation of events that are really generated in UTC. Ontologies are important; even within the FAA discussions, it was amusing to note the different uses of the term "surveillance" — air traffic controllers use it with respect to knowing where an airplane is; security folk would like to watch what's happening on the plane. Similarly, for an application so concerned with where things are, a common notion of and expression for geography and trajectory could prove critical.

If subscribers can retrieve events that happened before their subscription was entered, the system will need a mechanism for storing past events, matching them against new subscriptions, and forwarding the results. Storage issues also get tangled with legal requirements for data preservation and access.

An architecture based on bringing all events to a single locus for sorting against all subscriptions is unlikely to scale to the national airspace or provide the necessary reliability to actually fly

planes. Thus, a major issue in such a publish-and-subscribe system is the organization and replication of *brokers* (that is, event channels) that mediate between publishers and subscribers. Because the system must grow to encompass new applications, discovery and allocation are critical issues: matching brokers to the kinds of messages they receive and republish. We might find a need for a hierarchical broker topology, the ability for brokers to forward subscriptions (or parts of subscriptions) to each other, and a facility that lets a forwarding broker collate multiple answers back to a single form.

In airspace control, there are often several sources of the same information. For example, you can find out where an aircraft is by projecting from its original flight plan, reading a current radar trace, or extrapolating the aircraft's message about its GPS location and trajectory. The underlying architecture might not need to arbitrate between information sources. However, thinking about how to organize with respect to multiple information sources

IEEE INTERNET COMPUTING

IEEE Computer Society Publications Office
10662 Los Vaqueros Circle
Los Alamitos, CA 90720

EDITOR IN CHIEF

Robert E. Filman • filman@computer.org
ASSOCIATE EDITOR IN CHIEF
Li Gong • li.gong@sun.com

EDITORIAL BOARD

Jean Bacon • jean.bacon@cl.cam.ac.uk
Miroslav Benda • miro@amazon.com
Elisa Bertino • bertino@dsi.unimi.it
Scott Bradner • sob@harvard.edu
Siobhán Clarke • siobhan.clarke@cs.tcd.ie
Fred Douglass • f.douglass@computer.org
Stuart I. Feldman • sif@us.ibm.com
Ian Foster • foster@cs.uchicago.edu
Monika Henzinger • monika@google.com
Michael N. Huhns • huhns@sc.edu
Leonard Kleinrock • lk@cs.ucla.edu
Doug Lea • dl@altair.cs.oswego.edu
Frank Maurer • maurer@cpsc.ucalgary.ca
Daniel A. Menascé • menasce@cs.gmu.edu

Chris Metz • chmetz@cisco.com

Charles J. Petrie • petrie@nrc.stanford.edu
(EIC emeritus)

Krithi Ramamritham • krithi@cse.iitb.ac.in
Munindar P. Singh • singh@ncsu.edu
(EIC emeritus)

Craig Thompson • cwt@uark.edu
Steve Vinoski • vinoski@ieee.org
Jim Whitehead • ejw@soe.uscs.edu

IEEE Communications Society Liaison

G.S. Kuo • gskuo@ieee.nccu.edu.tw

STAFF

Lead Editor: Steve Woods
swoods@computer.org
Group Managing Editor: Gene Smarte
Staff Editors: Scott L. Andresen,
Kathy Clark-Fisher, and Jenny Ferrero
Production Editor: Monette Velasco
Magazine Assistant: Hazel Kosky
internet@computer.org
Graphic Artist: Alex Torres
Contributing Editors: David Clark, Greg Goth,
Keri Schreiner, Joan Taylor

Publisher: Angela Burgess

Assistant Publisher: Dick Price

Membership/Circulation Marketing

Manager: Georgann Carter

Business Development Manager:

Sandy Brown

Advertising Supervisor: Marian Anderson

CS Magazine Operations Committee

Bill Schilit (chair), Jean Bacon, Pradip Bose,
Doris L. Carver, George Cybenko, John C.
Dill, Frank E. Ferrante, Robert E. Filman,
Forouzan Golshani, David Alan Grier,
Rajesh Gupta, Warren Harrison,
Mahadev Satyanarayanan,
Nigel Shadbolt, Francis Sullivan

CS Publications Board

Michael R. Williams (chair), Michael Blaha,
Mark Christensen, Sorel Reisman,
Jon Rokne, Bill Schilit, Linda Shafer,
Steven L. Tanimoto, Anand Tripathi





is a critical architectural issue.

We have been discussing communication abstractions. Of course, building a real airspace control system involves concretely handling the issues of real systems:

- performing all the facets of security, including authentication, intrusion detection, and access control;
- providing reliability to a life-critical system;
- meeting the real-time and data streaming needs of applications;
- accounting for the consumption of resources;
- monitoring performance;
- enabling configuration management; and
- identifying faults.

Even if an ideal solution to all these problems emerged, Venus-like, whole and complete, an even larger issue is how to achieve a smooth and reliable transition from current practice. Airspace management is a working, life-critical, heterogeneous, real-time, always-running application used not only by large, well-identified client organizations (for example, airlines) but also by individuals (for example, private pilots in noncommercial aircraft). It can't be changed instantaneously; it can't be turned off to bring up a new version; it has to work; and there are many stakeholders with differing opinions about the value and desirability of improvements.

Airspace Information and the Internet

We already have a distributed information system in the Internet. How well do the Internet, TCP/IP, DNS, and their kin stack up against the requirements of airspace information management? A great virtue of the Internet is its arbitrary connectivity — any process can potentially communicate with any service. Arbitrary connectivity's importance for unexpected evolution cannot be exaggerated: traditional system

IC Welcomes New Editorial Board Member



Monika Henzinger is the research director at Google.

Prior to joining Google, she was a research staff member at Digital's Systems Research Center in Palo Alto, Calif., where she worked on Web information retrieval and systems performance measurements as well as graph algorithms and data structures, which are useful for

mapping the Web's topography and behavior.

Henzinger received a PhD from Princeton University in 1993. She then joined the Computer Science Department at Cornell University as an assistant professor and received a US National Science Foundation Career award in 1995. Henzinger is also a recipient of the Wallace Memorial Honoric Fellowship at Princeton University, a Studienstiftung des deutschen Volkes fellowship, and a Siemens Scholarship fellowship.

Net as a whole, the lack of inherent accounting mechanisms, and the limitations on network evolvability should be familiar to our readers. More interesting as an abstraction issue is the philosophical shift from addressed to content-directed communication. It will be interesting to see how radical a set of changes in our thinking about networks are needed to accomplish these goals.

Closing Remarks

The theme of the Communications Abstractions workshop at ECOOP in Oslo this June will be "communication abstractions for airspace management." Please submit a position paper if you'd like to come (<http://perso-info.enst-retagne.fr/~beugnard/ecoop/WS-CADS04-CFP.html>). My thanks to Steve Bradford, Brian Glass, Josh Hung, and Jack Levine for comments on the drafts of this column. My work on technologies for SWIM is supported by NASA's Airspace Program. □

development for applications like air traffic controls starts by drawing boxes connected by lines, and is flummoxed when new applications demand additional connections. It's also important to keep in mind the lessons of the end-to-end argument: put the intelligence in the applications, not the network. This universality must be balanced by the

knowledge that we are discussing a system for airspace information management, not unstructured communication.

That said, the current Internet architecture doesn't match our application in a variety of ways. The poor quality of IPv4 Internet security (which IPv6 might ameliorate), the lack of real-time capabilities, the difficulty in managing the

ADVERTISER / PRODUCT INDEX JANUARY / FEBRUARY 2004

Advertisers/Products	Page Number	Advertising Personnel	
CTIA Wireless 2004	Cover 4	Marion Delaney IEEE Media, Advertising Director Phone: +1 212 419 7766 Fax: +1 212 419 7589 Email: md.ieeemedia@ieee.org	Sandy Brown IEEE Computer Society, Business Development Manager Phone: +1 714 821 8380 Fax: +1 714 821 4010 Email: sb.ieeemedia@ieee.org
Email Security Summit 2004	Cover 2	Marian Anderson Advertising Coordinator Phone: +1 714 821 8380 Fax: +1 714 821 4010 Email: manderson@computer.org	
<i>Boldface denotes advertisements.</i>			
Advertising Sales Representatives			
Mid Atlantic (product/recruitment) Dawn Becker Phone: +1 732 772 0160 Fax: +1 732 772 0161 Email: db.ieeemedia@ieee.org	Midwest (product) Dave Jones Phone: +1 708 442 5633 Fax: +1 708 442 7620 Email: dj.ieeemedia@ieee.org	Midwest/Southwest (recruitment) Darcy Giovingo Phone: +1 847 498-4520 Fax: +1 847 498-5911 Email: dg.ieeemedia@ieee.org	Northwest/Southern CA (recruitment) Tim Matteson Phone: +1 310 836 4064 Fax: +1 310 836 4067 Email: tm.ieeemedia@ieee.org
New England (product) Jody Estabrook Phone: +1 978 244 0192 Fax: +1 978 244 0103 Email: je.ieeemedia@ieee.org	Will Hamilton Phone: +1 269 381 2156 Fax: +1 269 381 2556 Email: wh.ieeemedia@ieee.org	Southwest (product) Bill Wageneck Phone: +1 972 423 5507 Fax: +1 972 423 6858 Email: bill.wageneck@wageneckassociates.com	Japan German Tajiri Phone: +81 42 501 9551 Fax: +81 42 501 9552 Email: gt.ieeemedia@ieee.org
New England (recruitment) Barbara Lynch Phone: +1 401 739-7798 Fax: +1 401 739 7970 Email: bl.ieeemedia@ieee.org	Joe DiNardo Phone: +1 440 248 2456 Fax: +1 440 248 2594 Email: jd.ieeemedia@ieee.org	Northwest (product) Peter D. Scott Phone: +1 415 421-7950 Fax: +1 415 398-4156 Email: peterd@pscottassoc.com	Europe (product) Hilary Turnbull Phone: +44 1875 825700 Fax: +44 1875 825701 Email: impress@impressmedia.com
Connecticut (product) Stan Greenfield Phone: +1 203 938 2418 Fax: +1 203 938 3211 Email: greenco@optonline.net	Southeast (product/recruitment) C. William Bentz III Email: bb.ieeemedia@ieee.org Gregory Maddock Email: gm.ieeemedia@ieee.org Sarah K. Wiley Email: sh.ieeemedia@ieee.org Phone: +1 404 256 3800 Fax: +1 404 255 7942	Southern CA (product) Marshall Rubin Phone: +1 818 888 2407 Fax: +1 818 888 4907 Email: mr.ieeemedia@ieee.org	Europe (recruitment) Penny Lee Phone: +20 7405 7577 Fax: +20 7405 7506 Email: reception@essentialmedia.co.uk